

Monarc Simulation of ATLAS Data Processing

With 1/3 of all ESD data in every Tier1 center

September 20, 2001

Igor Baryakhtar^a, Krzysztof Śliwa^a, Marcin Wolter^{a,b}

^aTufts University, Boston

^bInstitute of Nuclear Physics, Krakow

Cluster Architecture

Tier0

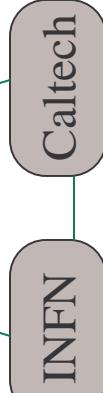
All centers are connected with each other by 200 MB/sec internet links



10000 MB/sec

RAW, ESD, AOD, TAG

Tier1



1/3 ESD, AOD, TAG in every center

Tier2



AOD, TAG in every center

Only few internet links are marked on the plot.

Processing Centers

TIER0

- 1200 nodes, 500 SI95 CPU each
- Internet connection - 200 MB/sec (10 000 MB/sec link between Tier0 and Tier1 at CERN)
- DB link speed 200 MB/sec

TIER1 (6 centers)

- 550 nodes, 500 SI95 CPU each
- Internet connection - 200 MB/sec (CERN-CERN_A 1000 MB/sec)
- DB link speed 200 MB/sec

TIER2 (4 centers)

- 30 nodes, 500 SI95 CPU each
- Internet connection - 200 MB/sec
- DB link speed 200 MB/sec

Data Processing at TIER0

RAW -> ESD:

- Processing: 2.7×10^9 events/year (115 days, 10^7 secs), i.e. 2.35×10^7 ev/day.
- Reprocessing: 1.55×10^9 events/(90 days) = 1.72×10^9 ev/day – reprocessing.
- Total RAW→ESD processing: 4.07×10^7 ev/day.
- RAW events: CPU: 640 SI95*sec/ev, Size 2 MB.
- **Data 81.4 TB/day , CPU $2.6 \times 10^{10} \text{ sec/day}$**

ESD -> AOD -> TAG

- Processing: 1.55×10^9 ev/year, i.e. 1.35×10^7 ev/day (less events due to selection)
- Reprocessing of RAW and consecutive ESD→AOD→TAG building: 1.55×10^9 ev/(90 days) = 1.72×10^7 ev/day.
- Independent reprocessing of ESD→AOD→TAG: 1.55×10^9 events/10 days
 $= 1.55 \times 10^8$ events/day.
- Total ESD→AOD→TAG processing: 1.86×10^8 events/day
- Size: ESD 0.5 MB, AOD 0.01MB, TAG 0.002MB
- CPU: ESD 50 SI95*sec, AOD 5.0 SI95*sec, TAG 0.25 SI95*sec.
- **Data 95 TB/day , CPU $1.02 \times 10^{10} \text{ sec/day}$**

Daily Data Processing at TIER1 Center

Selection (20 groups):

- Physics selection: 20 times all TAGs/day (in all six TIER1 centers together).

Selection analyze factors: RAW: 1/10000, ESD: 1/100, AOD: 1/10

Total data $254 \text{ TB/day (42 TB/day/center)}$, total CPU $40.73 * 10^9 \text{ sec/day (6.79 * 10^9 sec/day/center)}$

Analysis - in every Tier1:

- Physics analysis: 300 users^(a) (75% of all events)
Analyze factors: RAW: 1/10000, ESD: 1/100, AOD: 1
1/3 of ESD, all AODs and TAGs locally, RAWs at Tier0
- Data $20 \text{ TB/day/center, CPU } 6.76 * 10^9 \text{ sec/day/center}$**

(a) A user: analyses 0.25% of all events during one day

Daily Data Processing at TIER2 Center

Analysis - at every Tier2:

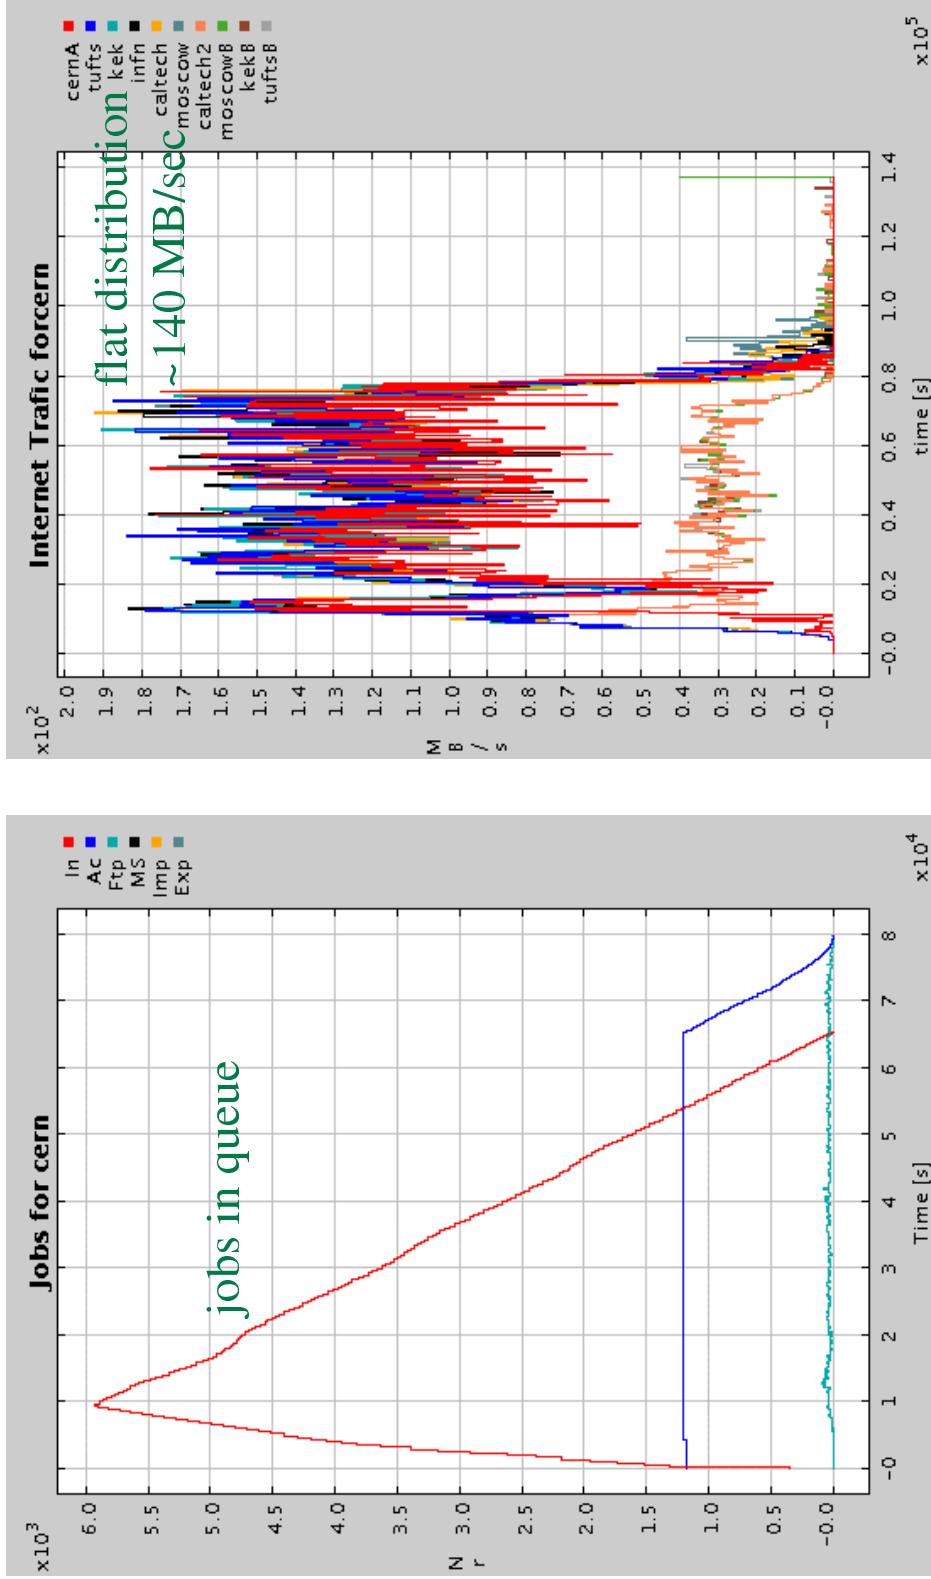
- Physics analysis: 60 users^(a) (15% of all events)

Analyze factors: RAW: 1/10000, ESD: 1/100, AOD: 1
AODs and TAGs locally, ESDs at Tier1, RAWs at Tier0

4 TB/day/center, CPU 1.35 sec/day/center

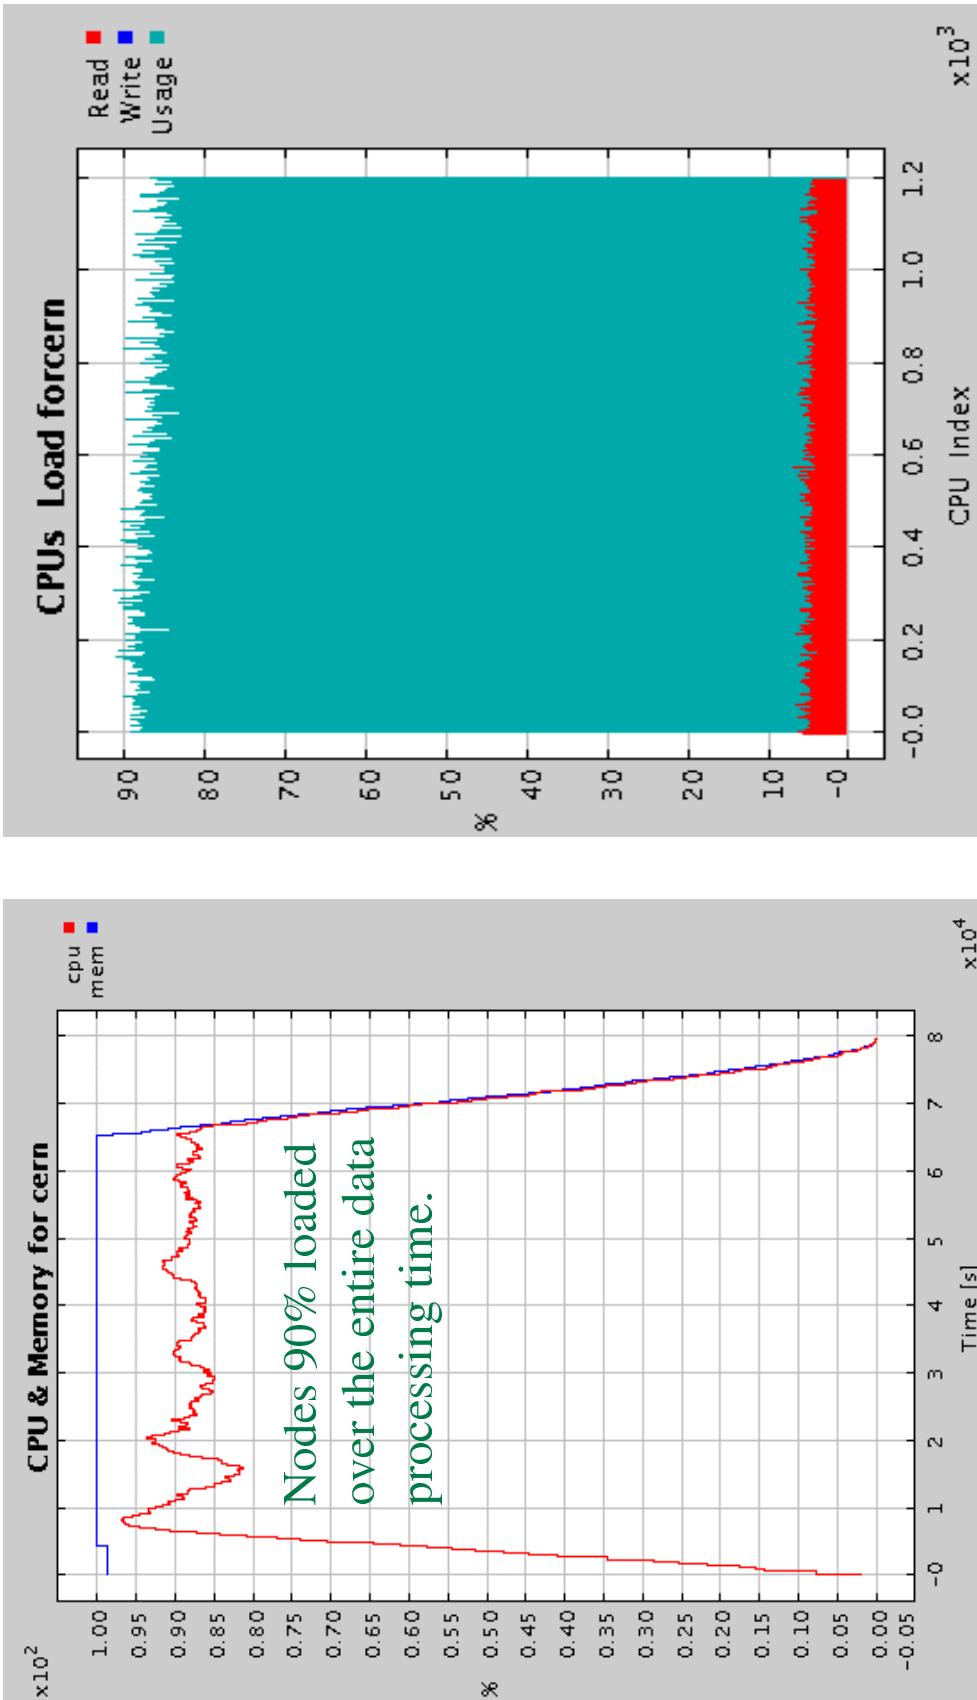
(a) A user: analyses 0.25% of all events during one day

Simulation Results – Tier0



Jobs and internet traffic at **TIER0**. Time needed to process and transfer data is approximately one day (one day = 8.6×10^4 sec).

Simulation Results – Tier0



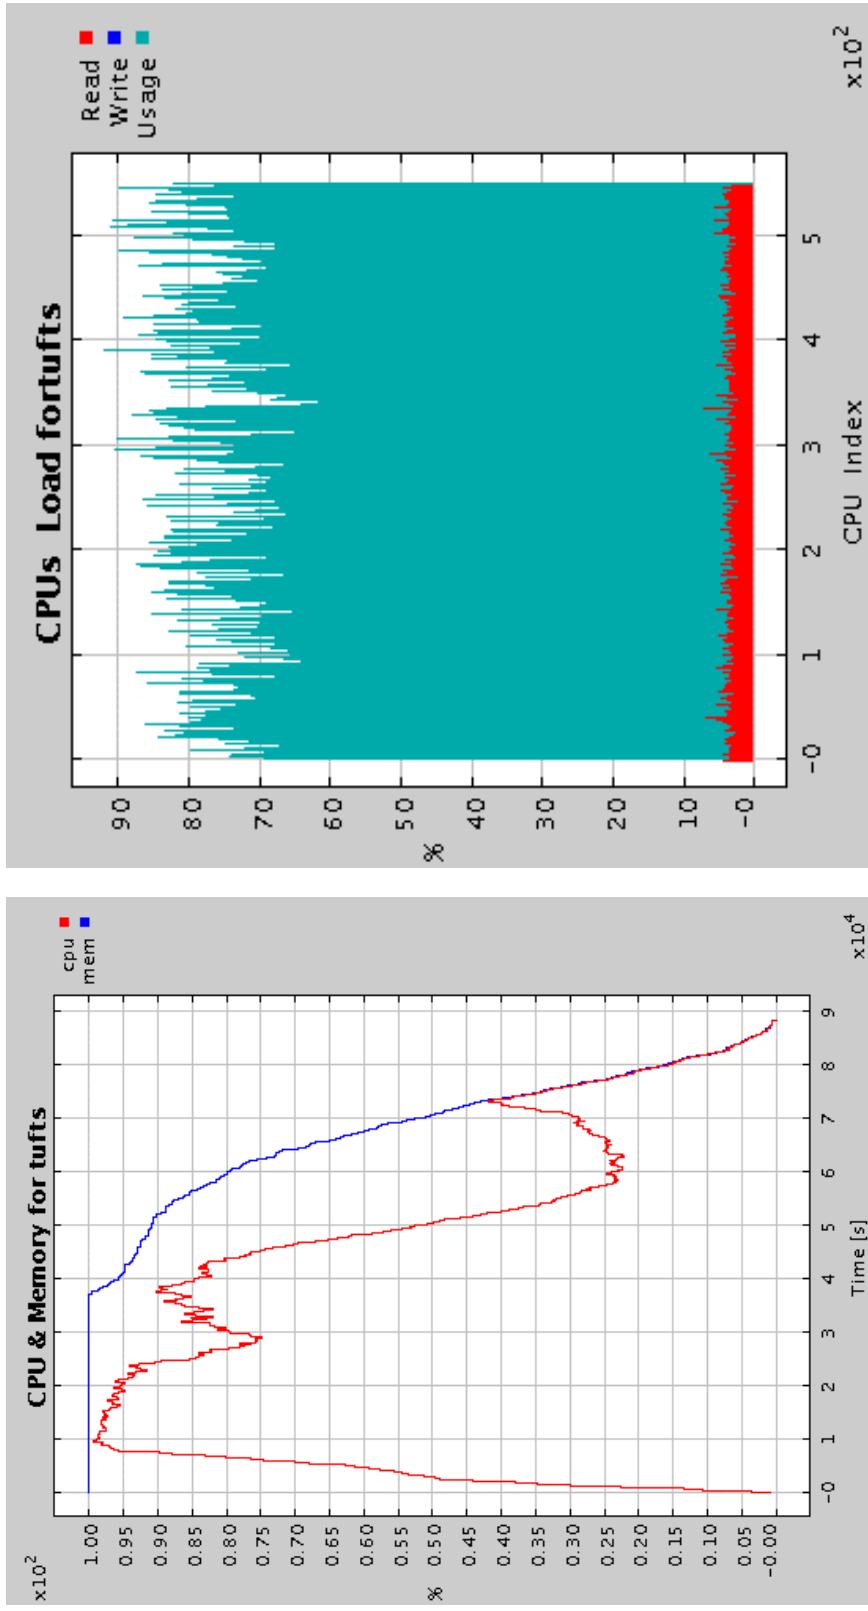
CPU/memory usage and node load at TIER0. Nodes are loaded in 90%.

20 September 2001

Monarc Simulation, Marcin Wolter

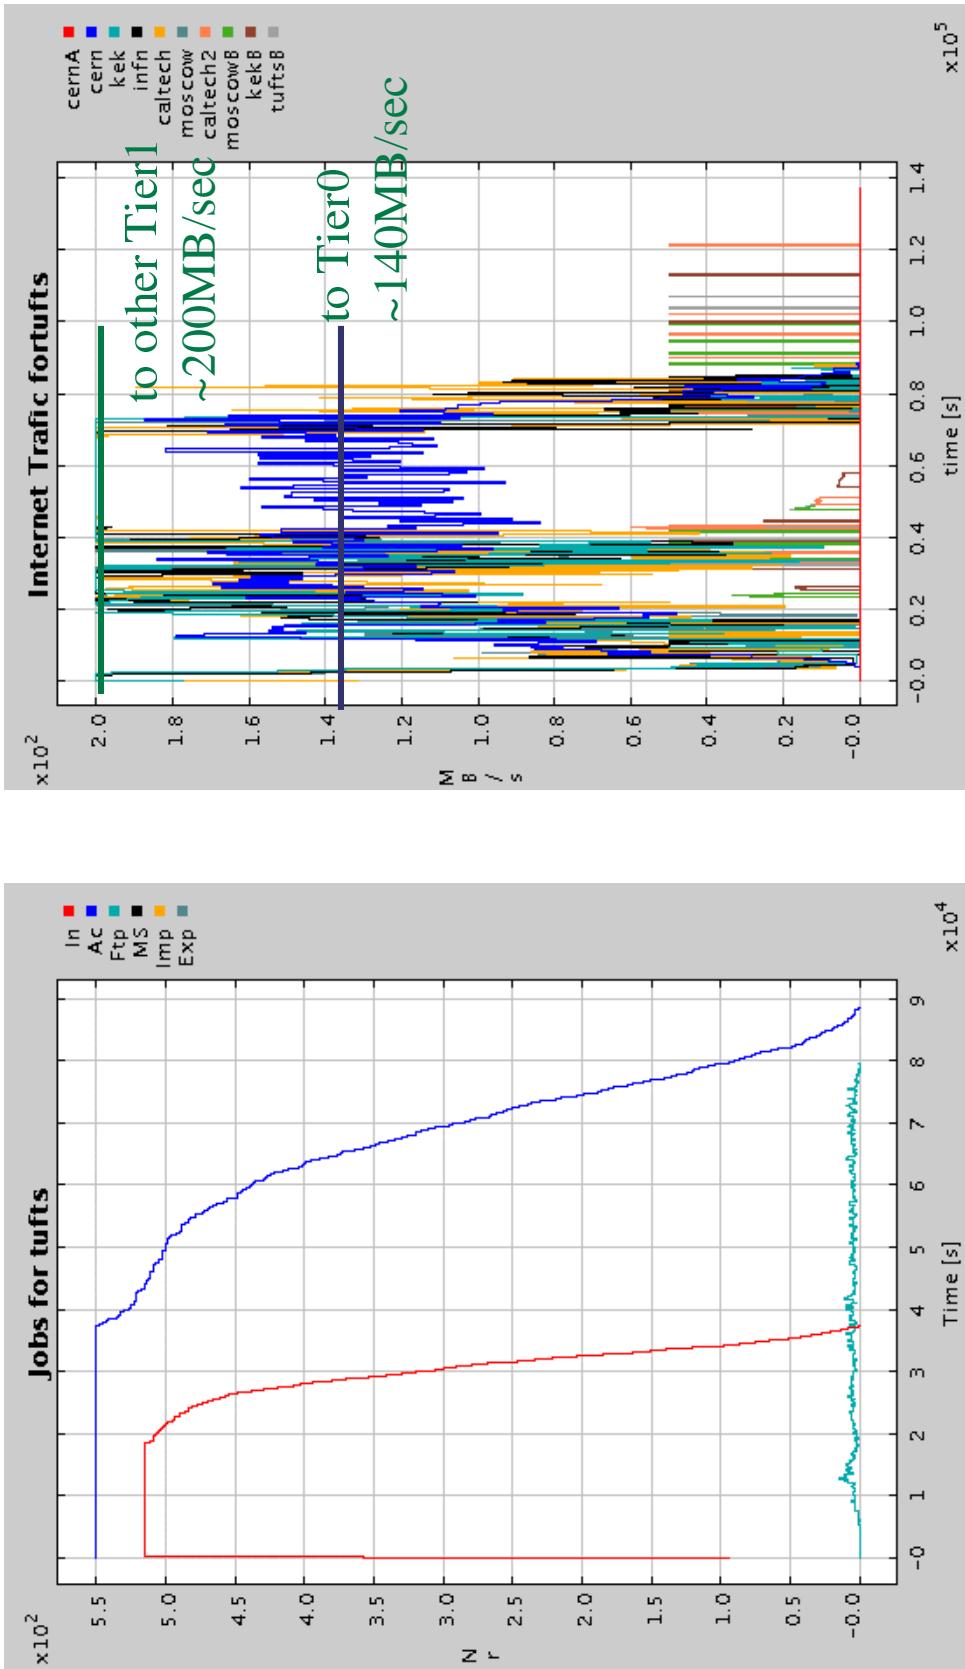
8

Simulation Results – TIER1



CPU/memory and node load at TIER1. Nodes are in average 75% loaded.
50% of CPU load is due to the selection, and 50% due to the analysis jobs.

Simulation Results – Tier1



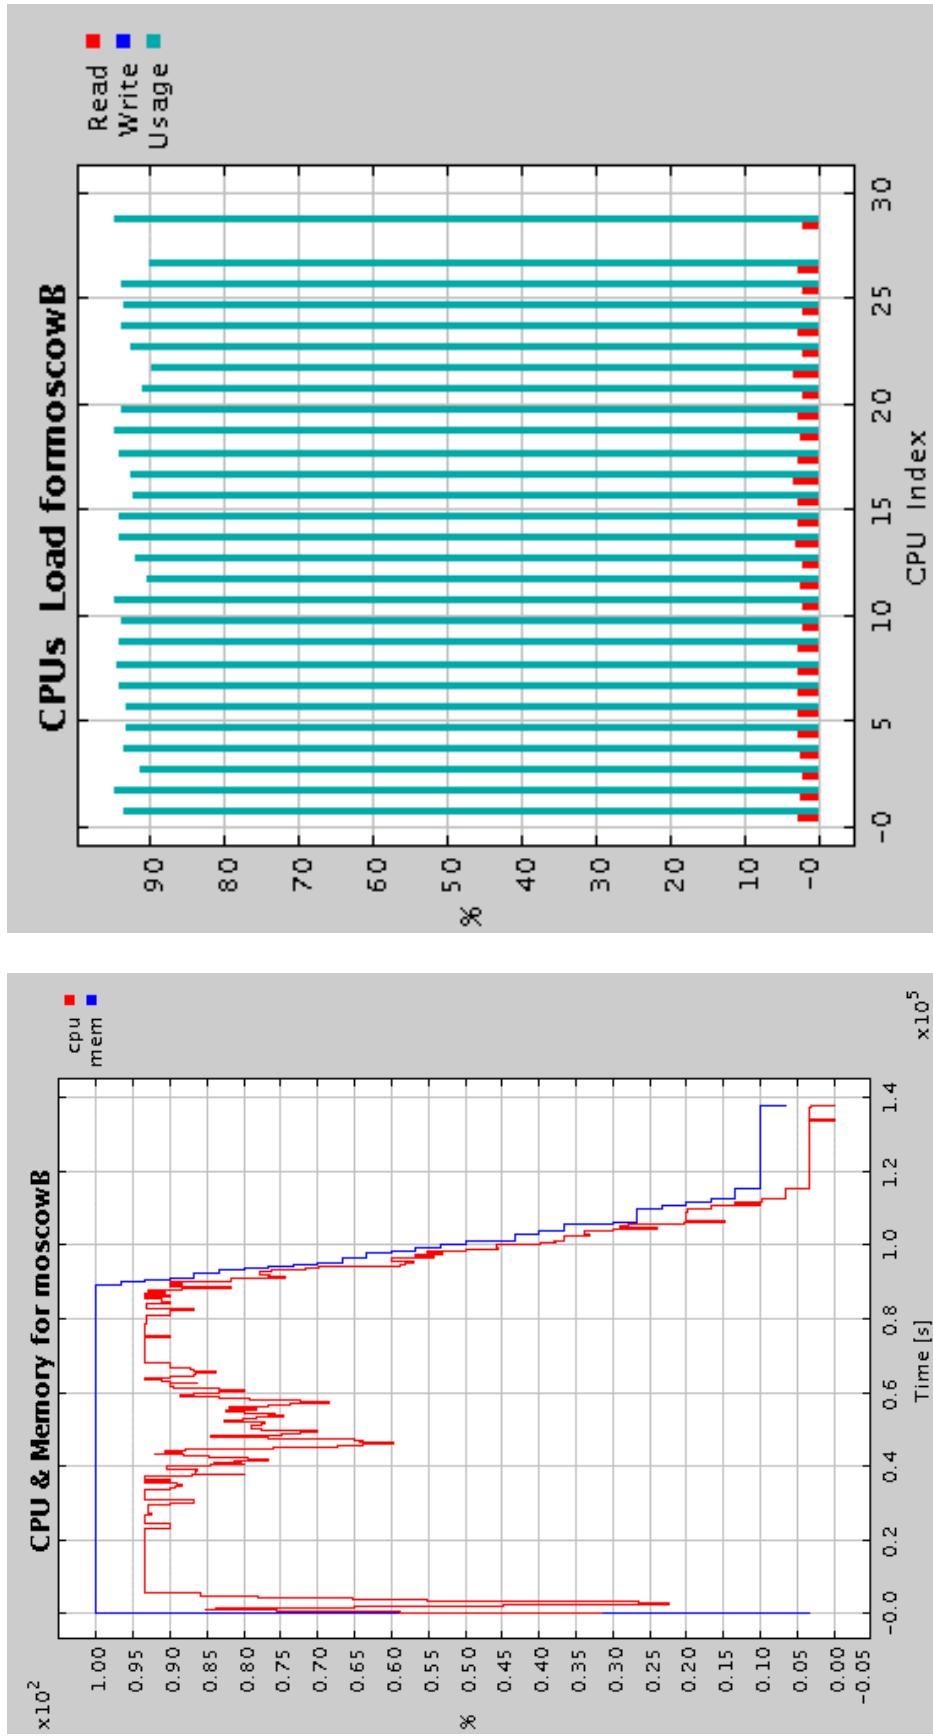
Jobs and internet traffic at TIER1.

20 September 2001

Monarc Simulation, Marcin Wolter

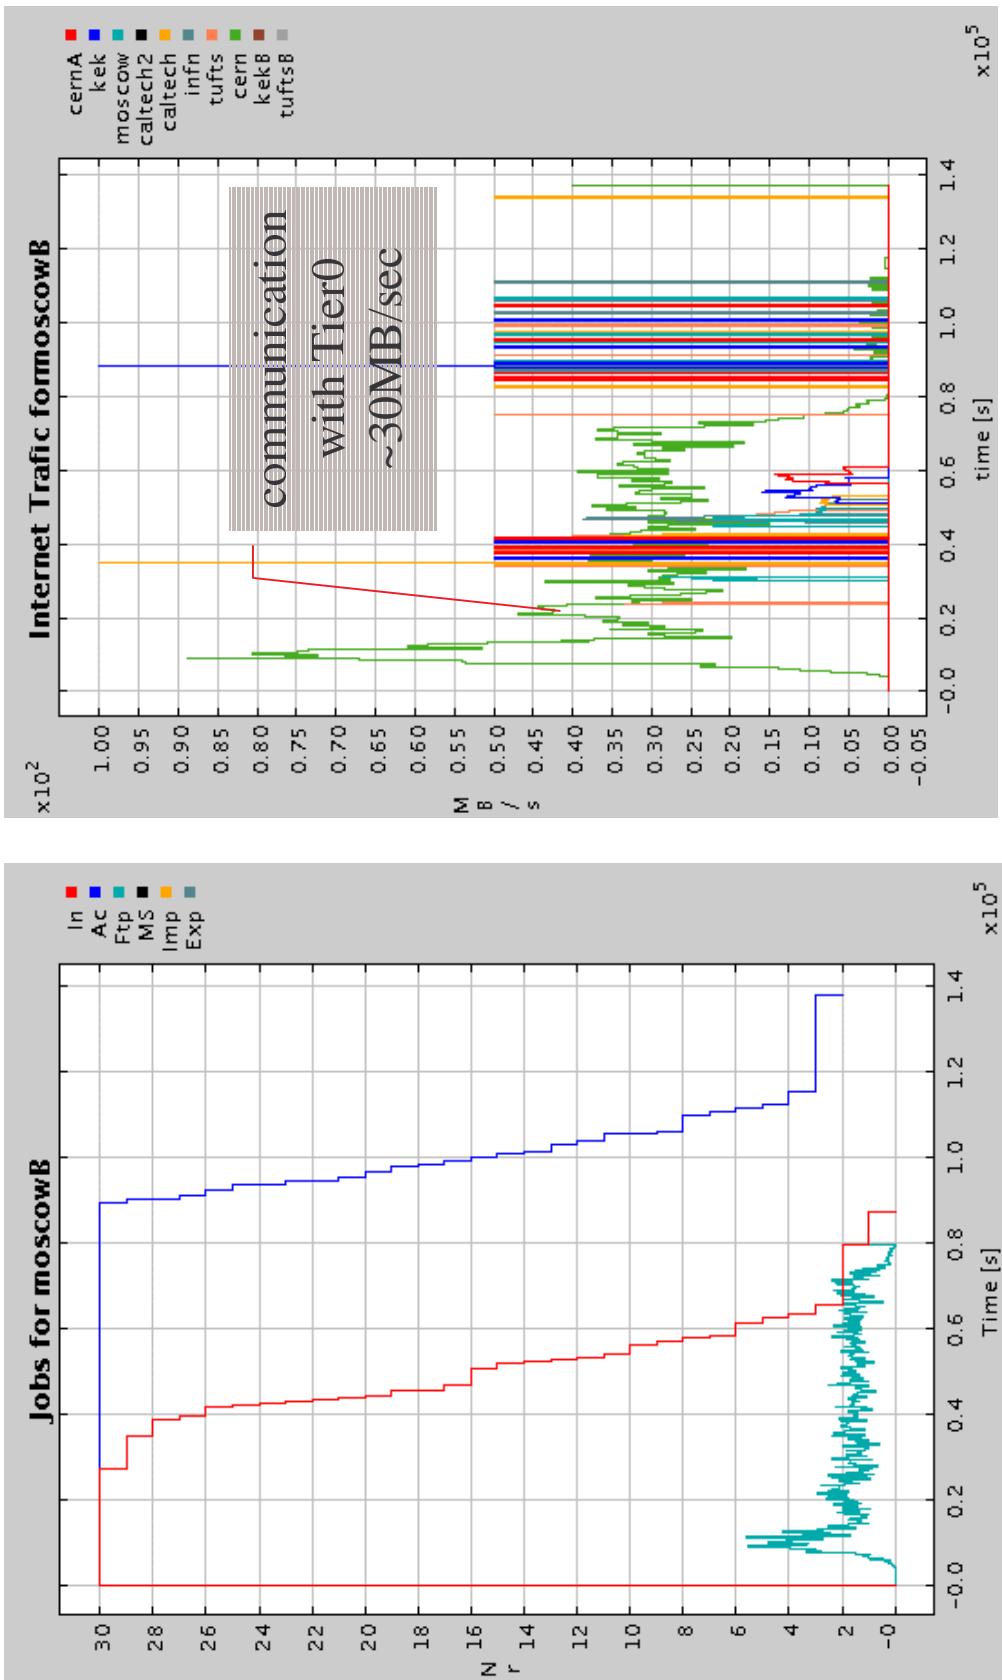
10

Simulation Results – TIER2



CPU/memory and node load at TIER2. Nodes are in average over 90% loaded.

Simulation Results – Tier2



Jobs and internet traffic at TIER2.

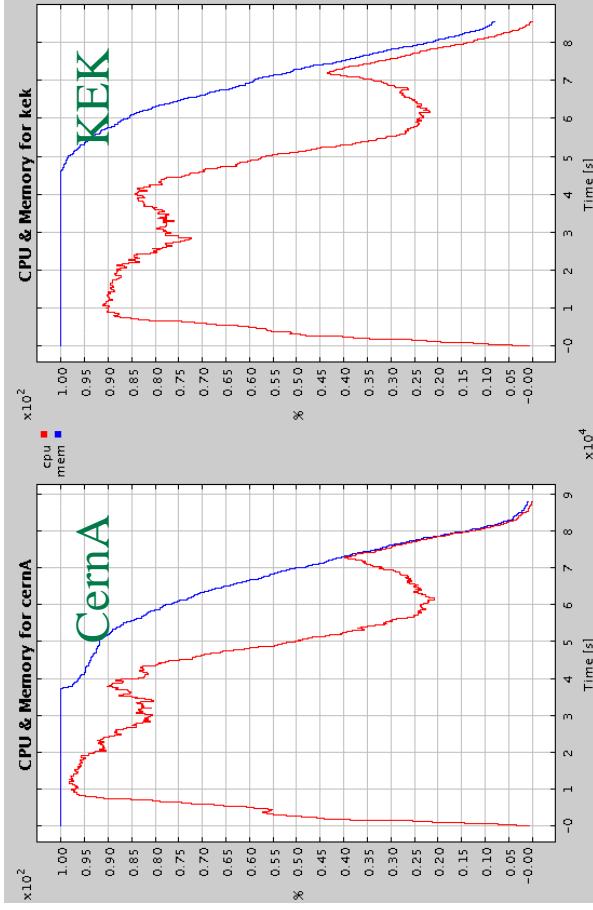
20 September 2001

Monarc Simulation, Marcin Wolter

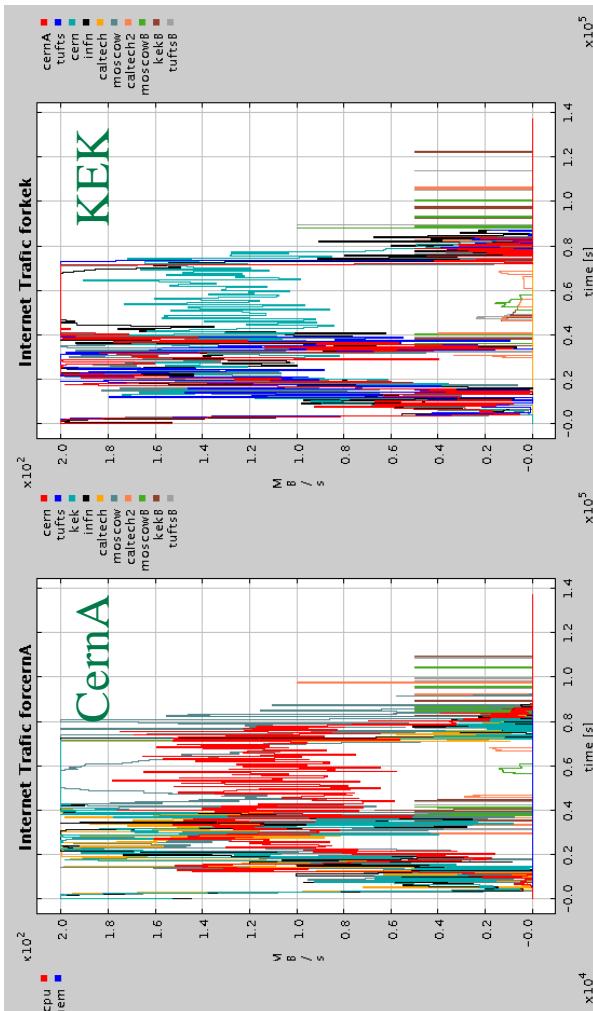
12

Comparison of CernA and another Tier1 center

CPU/memory usage



Internet traffic

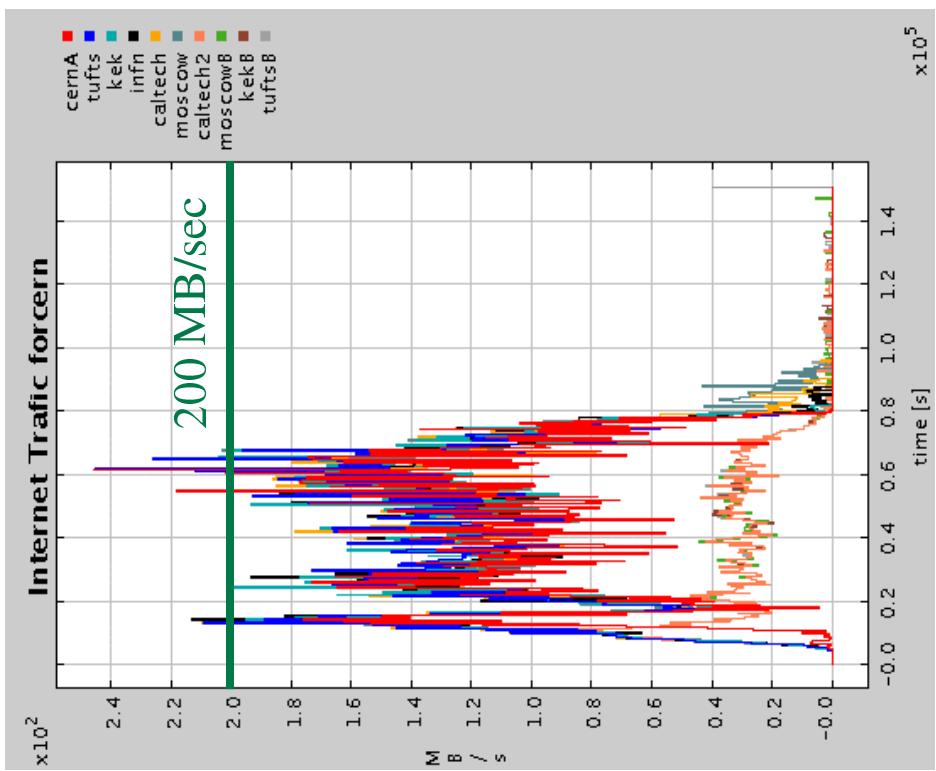


CERN has 50 times faster connection to CERN Tier0 than KEK.

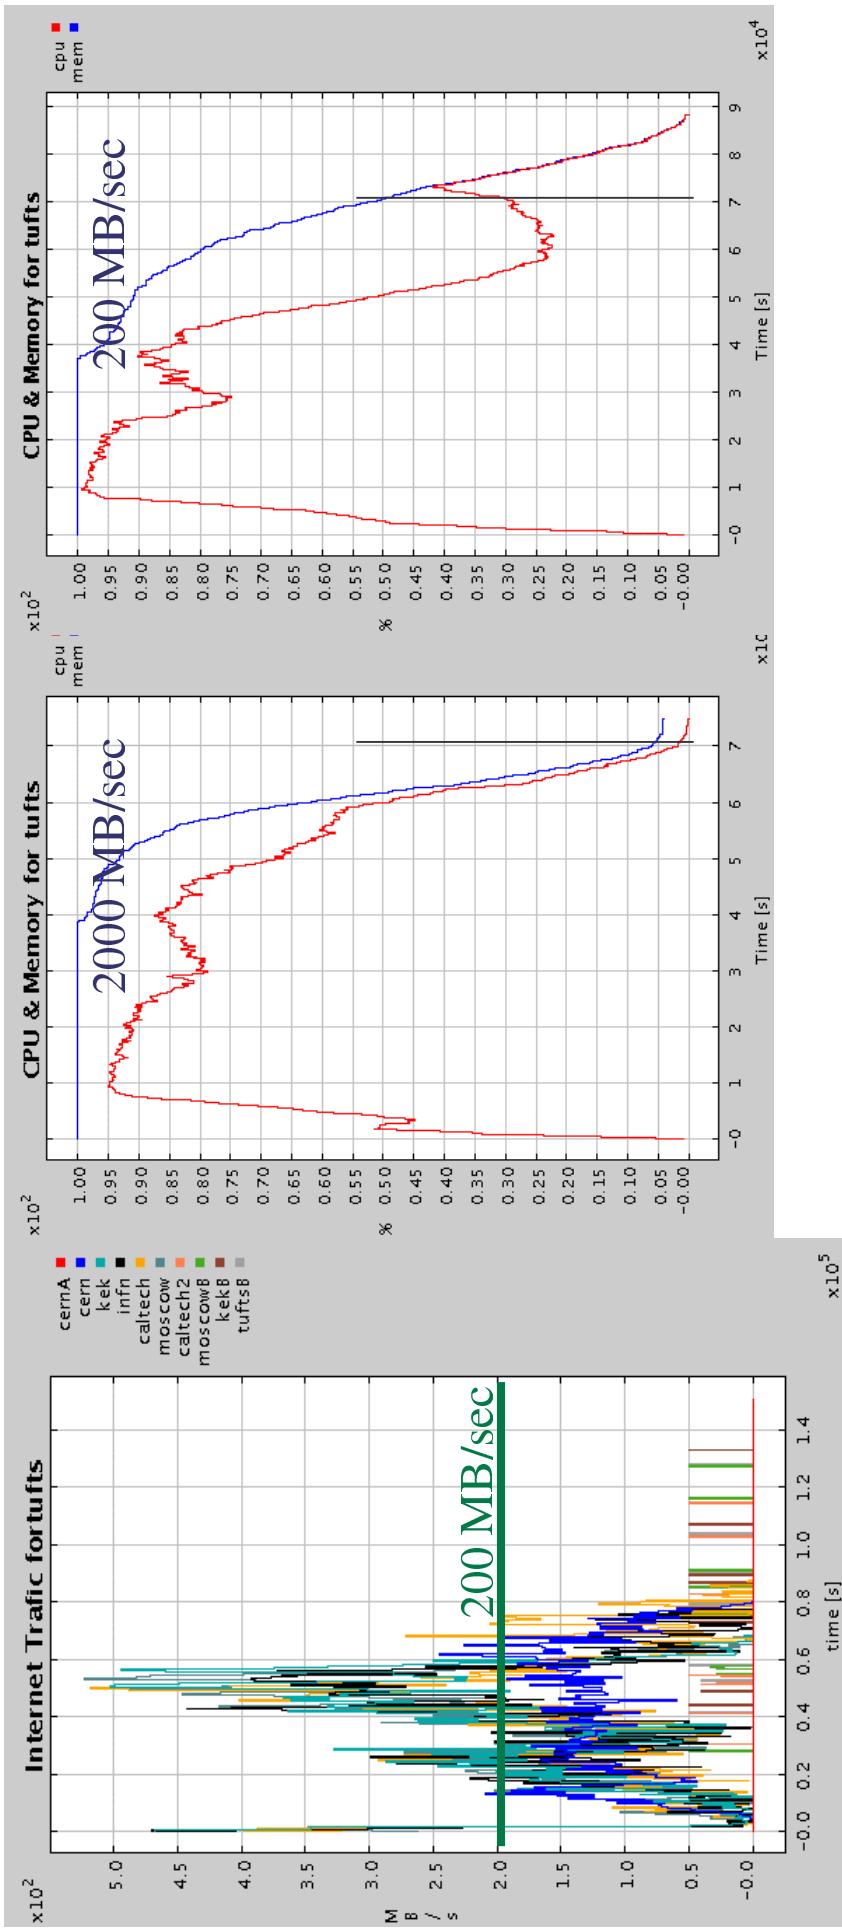
No significant difference observed.

High internet bandwidth 2000 MB/sec

Tier0 center never doesn't require an internet bandwidth higher than 200 MB/sec.



High internet bandwidth 2000 MB/sec



High internet bandwidth speeds up (~15%) the data processing at Tier1 centers (traffic between Tier1 centers, not to Tier0).

20 September 2001

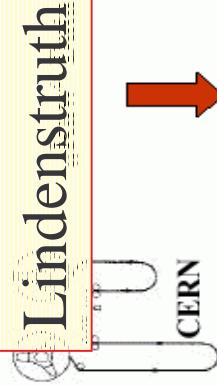
Monarc Simulation, Marcin Wolter

15

Conclusions

- ATLAS data processing can be performed with 1/3 of all ESD data stored in every Tier1 center.
- 200 MB/sec internet links are sufficient, however faster links between Tier1 centers slightly reduce the processing time.
- No faster link between Tier0 and Tier1's is needed.

Reference Data flow Architecture (per annum) for ATL from Volker



(1) Full Reconstruction
reprocessing

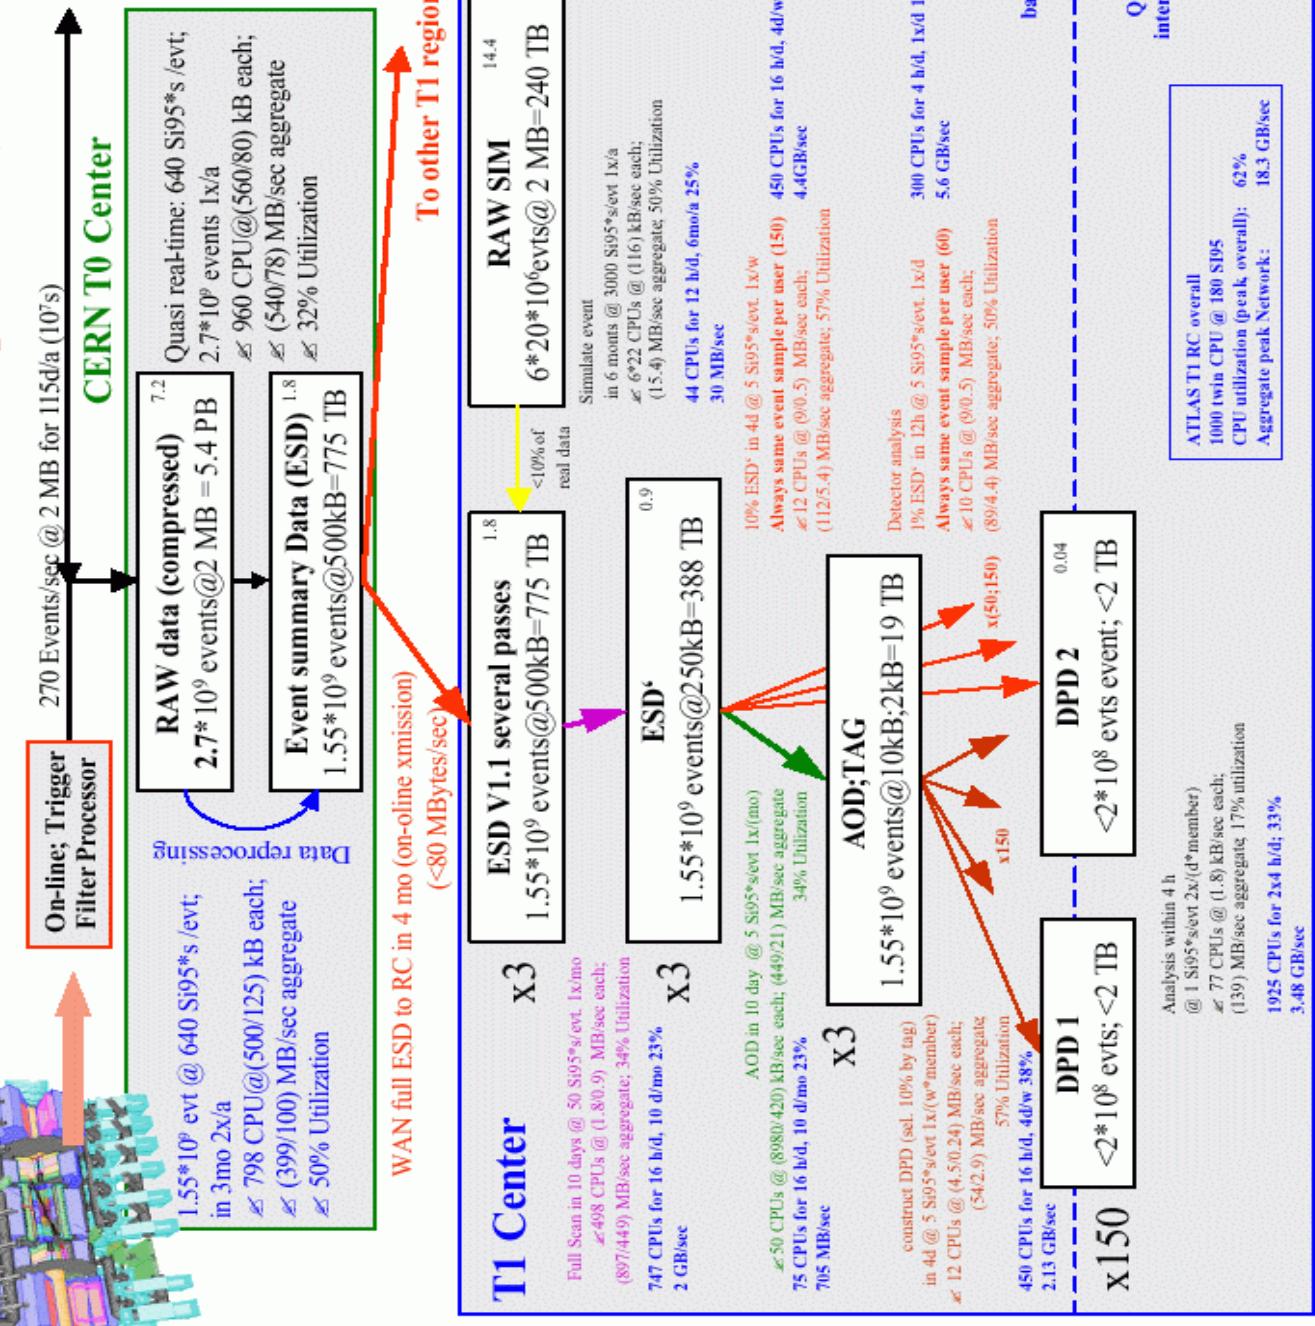
WAN full ESD to RC in 4 mo (on-line xmission)
(<80 MBbytes/sec)

(2) Construct physics ESD
(ESD')

(3) Construct AOD

(4) Construct n-tuple (DPD)

(5) Physics analysis



Tier1 CPU load by selection and analysis jobs

Analysis	TAG			
AOD		0.75*1*0.25		0.1875
ESD		0.75*1*5		3.7500
RAW		0.75*0.01*50		0.3750
Total:		0.75*0.0001*848		0.0636
			4.3761	
SELECTION	TAG			
AOD		1/6*20*1*0.25		0.8325
ESD		1/6*20*0.1*5		1.6650
RAW		1/6*20*0.01*50		1.6650
Total:		1/6*20*0.0001*848		0.2824
			4.4449	

(All numbers should be multiplied by a total number of events $1.55*10^9$)

~50% selection, ~50% analysis